

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing this collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number. PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

BEST AVAILABLE COPY

Availability Metrics for Frequency Management

Charles H. Jones, PhD
Edwards Air Force Base
charles.jones@edwards.af.mil

ABSTRACT

One approach to improving spectrum usage efficiency is to manage the scheduling of frequencies more effectively. The use of metrics to analyze frequency scheduling could aid frequency managers in a variety of ways. However, the basic question of what is a good metric for representing and analyzing spectral usage remains un-answered. Some metrics capture spectral occupancy. This paper introduces metrics that change the focus from occupancy to availability. Just because spectrum not in use does not mean it is available for use. A significant factor in creating unused but unusable spectrum is fragmentation. Mission profiles for spectrum usage can be considered rectangles in a time versus frequency grid. Even intelligent placement of these rectangles (i.e., the scheduling of several missions' spectrum usage) can not always utilize all portions of the spectrum. Availability metrics provide a way of numerically answering the question: What was the probability that another mission could have been scheduled? This is a much more practical question than: Did we occupy the entire spectrum? If another mission couldn't have been scheduled, then the entire spectrum was effectively used, even if the entire spectrum wasn't occupied.

INTRODUCTION

In its daily incarnation, frequency management is the process of scheduling blocks of frequencies and time in a non-interfering fashion. The unit of scheduling is a mission profile, which is a single contiguous block of frequencies over a contiguous period of time. Thus, geometrically, a mission profile can be considered a rectangle in a time vs. frequency grid. Figure 1 shows such a grid with four mission profiles scheduled. A fifth potential mission that conflicts with one of the already scheduled missions is also shown. This conflict is geometrically illustrated by the overlap of the rectangles.

A significant concern in frequency scheduling is that of spectrum fragmentation. (This is virtually identical to the concept of disk fragmentation on a computer.) The placement of mission profiles can easily lead to a fragmentation of the domain into pieces that are not usable. For example, the potential new mission in Figure 1 can not be scheduled anywhere without conflicting with the missions already scheduled. Within this example, a simple solution is to move the left-most mission down a notch or two and everything fits just fine. However, real world applications may have tens or hundreds of missions so that deciding what rectangles to move may not always be immediately obvious. In fact, many people spend all day trying to deconflict frequency schedules.

There are a variety of reasons why scheduling frequencies is hard. One reason is simply that scheduling is a mathematically difficult problem. This scheduling problem is an example of a dynamic bin packing problem. Technically, such problems are \mathcal{NP} -hard, meaning that they are exponentially difficult and, in the worst case, computers could take years or centuries to find

20040922006

optimal solutions. Other realities that make frequency scheduling difficult include: not all radios are tunable; missions often require multiple frequency ranges; changes require coordination, time, and often have cascading effects; frequency isn't the only thing that has to be rescheduled; prioritizations are not always well established; and a variety of other reasons.




Figure 1 – Time vs. Frequency Grid with scheduled and conflicting missions.¹

Because of the difficulties of scheduling, the reality is that the spectrum gets fragmented and not all of the spectrum is used. The questions addressed in this paper are: What metrics capture how well the scheduling was done considering the reality of fragmentation? How do we generate a numerical measure of whether or not another mission could have been scheduled? In other words, was there a *usable* portion of the spectrum *available*?

MISSION AVAILABILITY METRICS (FIXED TILE METHODS)

Focusing on the idea of whether a new mission profile will fit into a grid with already scheduled missions, one approach is simply to try every possible position and see if a conflict is created. That is, think of the mission profile as a fixed size, rectangular, tile and see if it fits into the jigsaw puzzle anywhere. Although, the question isn't just: Can a particular mission profile fit

¹ The mission profiles in all examples are fictional. In particular, most of the profiles are larger than in practice. They are provided for illustration of the concepts only.

somewhere? The question is more about quantifying probabilities of scheduling missions in general. So a next level use of the fixed tile approach is to ask: How many places can a mission profile be placed? This allows the calculation of a probability. The following example will illustrate this concept. In order to simplify the example, a time restriction is placed on the scheduling request. That is, the mission must be scheduled at a specific time, but the frequency is variable. This particular example is thus of the Time Required Mission Availability (TRMA) metric. Other metric variations will be outlined after the example. The formal algorithm for TRMA will be stated first and then an example calculation will be provided.

The algorithm requires these predetermined inputs:

1. Start time
2. Required bandwidth
3. Required duration
4. Available frequency range
5. Existing scheduled missions

Algorithm

Calculate lowest and highest frequency from available frequency range and required bandwidth.

```
available count = 0
for frequency = lowest frequency to highest frequency step dB
    if schedulable then
        available count = available count + 1
    end if
end for

number of available frequencies = (highest frequency - lowest
frequency) / dB

TRMA = available count / number of available frequencies
```

Example

Given the following inputs to the algorithm we can calculate TRMA.

- | | |
|--------------------------------|----------------|
| 1. Start time | 0600 |
| 2. Required bandwidth | 15 MHz |
| 3. Required duration | 5 hrs |
| 4. Available frequency range | 2200-2290 MHz |
| 5. Existing scheduled missions | (See Figure 2) |

Lowest Frequency = 2200

Highest Frequency = 2290 - 15 = 2275

Number of Available Frequencies = (2275-2200)/5 = 15

Available Count = 3 (Can be scheduled at 2200, 2205, and 2210 MHz)

TRMA = 3/15 = 0.2 (or 20 percent)

Example Discussion

TRMA is the probability of scheduling a mission given a mission profile and a required start time, but not a required frequency. Thus, for this example, there is a 20 percent probability of scheduling this mission. The converse interpretation is that, for this particular mission need, 80 percent of the spectrum is unusable. Contrast this to spectral occupancy. For the same scenario illustrated in Figure 2, for the time period 0600 to 1100 and frequency range 2200-2290, the Percent of Occupancy (PO) = $(2*5 + 2*1 + 3*7) / (5*18) = 31 / 90 = 0.34$ (or 34 percent). This gives the impression that there is a 66 percent possibility of scheduling this mission in contrast to an availability of 20 percent. Further, consider that if the required bandwidth were 30 MHz then TRMA would be 0 percent meaning that the mission could not be scheduled due to fragmentation even with a low PO.




Figure 2 – Example for Time Required Mission Availability (TRMA)

In general, numeric interpretations of TRMA include:

1. TRMA > 0 means the mission can be scheduled at that start time.
2. TRMA = 1 means there are no other missions scheduled at that start time.
3. The greater TRMA is, the more flexibility there is to schedule the mission.

Similar interpretations can be made for other variant availability metrics.

Other Variations

Using the basic concept of the fixed tile method, several different metrics can be defined.²

1. Frequency Required Mission Availability (FRMA). The probability of scheduling a mission given a mission profile and a required frequency, but not a required start time.
2. Ad Hoc Mission Availability (AHMA). The probability of scheduling a mission given a mission profile and flexibility in both frequency and start time.
3. Average Typical Mission Availability (ATMA). Most users doing frequency scheduling have more than one typical mission profile. ATMA is the average of the AHMA for all typical mission profiles. (Although, certainly you could average any of these metrics.)
4. Percent Fragmentation. Define a portion of the frequency-time domain to be fragmented if that portion is not usable to schedule a typical mission profile. Percent fragmentation can then be calculated as the percent unusable out of the total domain or as the percent unusable out of the portion of the domain not already scheduled. Since this is defined for a specific mission profile, the final percent fragmentation is defined to be the average fragmentation for all typical mission profiles.

The introduction of typical mission profiles allows individual users to customize some of the metrics to their particular environment. Another customizable aspect for all the metrics is the definition of both the total frequency available and the portion of the day used for the calculations.

MAXIMUM AVAILABILITY METRICS (VARIABLE TILE METHODS)

Instead of asking whether a particular mission profile can be scheduled it might be desirable to ask how large a profile can be scheduled. This question can be answered by using a similar method of looking at all possibilities, but this time, instead of using fixed-size tiles, the tiles are varied in size until they intersect existing scheduled missions. For example, to find the largest possible mission profile that can be scheduled at time 0:00 and 2290 MHz, start by placing a 1x1 tile in that location. If that doesn't intersect a scheduled mission, then try a 1x2 tile, and then a 1x3 tile, etc. until the tile does intersect. At that point, say at a 1x19 sized tile, then start increasing in the X direction and try a 2x19 tile. With a carefully crafted algorithm that checks all possible tiles, the largest such tile (or mission profile) can be found. In the example illustrated by Figure 2, the largest tile that can be placed at time 0:00 and 2290 MHz is a 7x6 tile representing a mission profile of 7 hours by 30 MHz. It is the largest mission profile in the sense of largest mission occupancy, which is calculated by multiplying the time by the bandwidth. In this case, the mission occupancy is 210 MHz hrs.

Just as with the fixed tile methods, there are several variant metrics that can be defined.

1. Maximum Available Duration (MAD). The longest possible duration a mission can be scheduled for a given start time and required bandwidth.
2. Maximum Available Bandwidth (MAB). The largest possible bandwidth a mission can be scheduled with for a given start time and required duration.
3. Maximum Available Mission Occupancy (MAMO). The mission profile with the largest mission occupancy that can be scheduled for a given start time and frequency.
4. Maximum MAMO. The maximum MAMO over a set of start times and frequencies.

² These and the maximum availability metrics listed have all been formally defined by the author.

5. Averages for each of these over a set of start times and frequencies can also be calculated.

The above example derives the MAMO for start time 0:00 and 2290 MHz of 210 MHz hrs. As with the fixed tile methods, some general numeric interpretations can be made. For example:

1. A large MAMO indicates a small amount of fragmentation and that a mission with occupancy less than the MAMO can probably be scheduled.
2. A small MAMO indicates that almost no new missions can be scheduled.

Some of these metrics lend themselves to 3-dimensional representation. For example, Figure 3 shows the MADs for the example illustrated in Figure 2. Similarly, Figure 4 shows the MAMOs for the example illustrated in Figure 2.




Figure 3 – Maximum Available Duration (MAD) Over Time and Frequency

With a little thought, there are some general interpretations of these graphs that can be made. The flat slope along the back portion of Figure 3 represents the fact that the grid used contains a finite amount of time. Thus for a 5 MHz bandwidth starting at 0:00 there is an available duration of 14 hours. But for every hour later the mission starts, the available duration decreases by an hour. Further, it is obvious that large bandwidths, even for a short duration, are virtually impossible to schedule. Perhaps the most immediate thing to take away from the MAMO chart in Figure 4 is how much potential mission occupancy is lost with very few missions scheduled and a very small percent of occupancy. The percent of the grid in Figure 2 occupied by scheduled missions is 19 percent. Consider that, with no missions scheduled the MAMO graph would be a flat plane going through the points (0:00, 2290 MHz, 1260 MAMO), (0:00, 2200 MHz, 70 MAMO), (13:00, 2290 MHz, 0 MAMO), and (13:00, 2200 MHz, 0 MAMO). That is, the maximum MAMO is 1260. Whereas, the scale for the MAMO in Figure 4 only goes up to 250 and has deep valleys in it! A little bit of fragmentation wipes out potential mission occupancy dramatically.




Figure 4 – Maximum Available Mission Occupancies (MAMO)

USE OF THESE METRICS

There are always many ways to use these and other metrics. Here are some suggestions.

Availability Trend Analysis

Many of these metrics could be used for long term planning – in both time and frequency domains. For example, if average AHMA for evenings or weekends is extremely high while average AHMA during normal workdays is nearly 0, then the metrics would indicate the need for more flights scheduled during evenings and weekends. Similarly, comparisons of AHMA for L and S bands could indicate the need to migrate one way or the other.

Another use would be in analyzing future ability to support a large program. For example, consider the situation where a new program appears. It is likely they will be able to provide an estimate of spectrum requirements over the life of the project. There will be existing projects that will continue to operate over all or part of the time period of this new project. It would be possible to take the requests that these existing projects have made and forecast their future usage. Add these to the new project's estimates. Run scheduling algorithms over future time periods and see what kinds of availability metrics are returned. If general availability is low, there should be real concern that the new project could not be supported and that further efficiency efforts should be implemented.

Defense Against Further Spectrum Reduction

This is certainly one of the driving forces for introducing these metrics. Actual percent of occupancy has never been, and probably never will be, 100 percent for Test and Evaluation (T&E) purposes (except possibly for short periods of time.) This makes defending usage

difficult. However, this doesn't change the fact that when the T&E community needs spectrum, either the spectrum is available or the mission doesn't happen. A proactive approach to spectrum defense is to present spectral usage data in a manner that captures the full complexity of the problem. Availability metrics are an attempt to do so.

As has been shown in the examples, availability may be very small even with a very small percent of occupancy. Although actual implementation of these metrics will provide the real answer, it is expected that most of the availability metrics are close to zero over normal operational times. This should help illustrate numerically that the T&E community is not wasting spectrum; that the spectrum is, in fact, being used efficiently within the working constraints of the discipline.

Real Time Reassignment Aids

To the person on the front line trying to schedule missions, the mission availability metrics could be useful, especially if the specific time and frequency choices were returned in addition to the counts. That is, instead of having to manually search for possible positions to schedule a requested mission, it would be quite useful if the schedulers could be provided with all possible choices. The maximum availability metrics could also be helpful, especially for ad hoc requests for projects that are willing to take what they can get.

Scheduling Algorithm Analysis

This may be wishful thinking, but it would be nice if the scheduling process could be automated or, at least, semi-automated. That is, are there algorithms that can help increase scheduling efficiency in both manpower and number of missions scheduled? Several of the metrics identified could be used for static analysis. That is, given a set of either past requests, fictional requests, or projected requests, it would be possible to run several algorithms and compare the resulting metrics.

DISCUSSION

To a certain extent, these metrics have been developed by starting with a couple of algorithmic ideas and defining the various metrics that come out of the process. However, these metrics also have been developed starting from the idea that, just because a portion of the spectrum is not occupied, doesn't mean that it is not used efficiently. A phrase that helps codify this is: "Use" is "denial to others." A simple example of this is that when you schedule a block of frequencies, no one else can use it whether you do or not. A more complicated example is fragmentation. If, all constraints considered, a fragmented but unused portion of the spectrum can not be used by a particular mission, that mission has been denied use of that portion of the spectrum. Hopefully, these ideas illustrate that a metric based solely on spectral occupancy is a somewhat uninformed metric.

Availability metrics are not currently in use. There is some potential for them to be incorporated into the Integrated Frequency Deconfliction System (IFDS), but a final decision has not been made. The frequency management community is engaged in a discussion of how to present spectral usage to the world. There are certainly other metrics or variants to be considered, but the author believes that the idea of availability, rather than occupancy, is a valid and strong candidate for capturing actual scheduling efficiency.



Availability Metrics for Frequency Management

Charles H. Jones, PhD
Edwards Air Force Base
charles.jones@edwards.af.mil



Outline

- Introduction
- Mission Availability Metrics
 - Fixed Tile Methods
 - Example of Time Required Mission Availability
- Maximum Availability Metrics
 - Variable Tile Methods
 - 3-D Representations
- Use of These Metrics
- Discussion




Frequency Management

- On daily basis, frequency management is scheduling blocks of frequencies and time in a noninterfering fashion.
- Each scheduled item is a mission profile.
- Geometrically, mission profiles can be thought of as rectangles in a time vs. frequency grid.
- Better scheduling can lead to more efficient spectrum use.
- Better metrics can lead to better scheduling.



Time vs. Frequency Grid With Scheduled Missions





Scheduling is Hard

- Technically, this type of scheduling is \mathcal{NP} -hard.
 - Exponentially difficult.
 - Computers might take years to find solutions.
- Fragmentation (similar to disk fragmentation) is a major concern.
 - The potential new mission in previous slide can not be scheduled anywhere.
- Many other practical reasons why this is hard.
 - Multiple frequencies per mission.
 - Ripple affects of changes.
 - Non-tunable radios.



Mission Availability Metrics (Fixed Tile Methods)

- Could another mission have been scheduled?
 - Was there *usable* spectrum *available*?
 - Much different from: Was the spectrum *occupied*?
- Check to see if a given mission would fit.
 - Use fixed rectangular tile (mission profile).
 - Systematically check all possible positions.
- Example: Time Required Mission Availability (TRMA)
 - No flexibility in time.
 - Flexibility in frequency.
 - Given mission profile.
 - Check all positions for the required time.



TRMA Algorithm

Calculate lowest and highest frequency from available frequency range and required bandwidth.

available count = 0

for frequency = lowest freq to highest freq step dB

 if schedulable then

 available count = available count + 1

 end if


end for

number of available frequencies = (highest frequency – lowest frequency) / dB

TRMA = available count / number of available frequencies



Example for Time Required Mission Availability (TRMA)





Numeric Example for TRMA



- Given the following inputs to the algorithm we can calculate TRMA.
 - Start time 0600
 - Required bandwidth 15 MHz
 - Required duration 5 hrs
 - Available frequency range 2200-2290 MHz
 - Existing scheduled missions (See Figure)



Numeric Example for TRMA (concl.)



- Lowest Frequency = 2200
- Highest Frequency = $2290 - 15 = 2275$
- Number of Available Frequencies = $(2275 - 2200)/5 = 15$
- Available Count = 3 (Can be scheduled at 2200, 2205, and 2210 MHz.)
- TRMA = $3/15 = 0.2$ (or 20%)



TRMA vs. Occupancy

- For the given time in the example:
 - Spectral Percent Occupancy = 34%.
 - Gives impression of 66% chance of scheduling the mission.
- In contrast:
 - TRMA says there is only 20% chance of scheduling the mission!
 - Implies 80% of the spectrum is unavailable and unusable for that mission!



Other Availability Metrics

- Frequency Required Mission Availability (FRMA).
 - The probability of scheduling a mission given a mission profile and a required frequency, but not a required start time.
- Ad Hoc Mission Availability (AHMA).
 - The probability of scheduling a mission given a mission profile and flexibility in both frequency and start time.
- Average Typical Mission Availability (ATMA).
 - Define a set of typical mission profiles.
 - The average of the AHMA for all typical mission profiles.
- Percent Fragmentation.
 - Define a portion of the frequency-time domain to be fragmented if that portion is not usable to schedule a typical mission profile.




Maximum Availability Metrics (Variable Tile Methods)

- What's the largest possible mission profile that can be scheduled?
 - Same systematic approach but vary the size of the tile.
- Maximum Available Duration (MAD).
 - The longest possible duration a mission can be scheduled for a given start time and required bandwidth.
- Maximum Available Bandwidth (MAB).
 - The largest possible bandwidth a mission can be scheduled with for a given start time and required duration.
- Maximum Available Mission Occupancy (MAMO).
 - The mission profile with the largest mission occupancy that can be scheduled for a given start time and frequency



Maximum Available Duration (MAD) Over Time and Frequency






3D Max Available Duration (MAD) Notes



- Back portion of graph has flat slope.
 - Artifact of finite time domain.
 - 14 hours at time 0:00.
 - Decrease 1 hour for every hour later start.
- Obvious that large bandwidths can not be scheduled.
- Large portions of availability lost with only:
 - 4 missions scheduled.
 - 19% occupancy.



Maximum Available Mission Occupancies (MAMO)





3D Max Available Mission Occupancy (MAMO) Notes



- With no missions scheduled:
 - Flat plane through points:
 - (0:00, 2290 MHz, 1260 MAMO),
 - (0:00, 2200 MHz, 70 MAMO),
 - (13:00, 2290 MHz, 0 MAMO), and
 - (13:00, 2200 MHz, 0 MAMO).
 - Maximum MAMO is 1260.
- 4 missions scheduled and 19% occupancy:
 - Maximum MAMO less than 250!
 - Large valleys in graph.



Use Of These Metrics



- Availability trend analysis.
- Defense against further spectrum reduction.
 - Occupancy unlikely to be 100%-in T&E.
 - Present usage data in a way that captures real complexities of scheduling.
- Real time reassignment aids.
- Scheduling algorithm analysis.



Discussion

- Metrics partially developed through combinatorial analysis of methods.
- Also developed from idea of "Use" as "Denial to Others."
 - Scheduled means you can't use it even if the people who scheduled it don't either.
 - Fragmentation denies use.
- Let's be proactive in presenting usage data.
- Suggest incorporating these metrics into the Integrated Frequency Deconfliction System (IFDS).